On a Network Creation Game

Joint work with Ankur Luthra, Elitza Maneva, Christos H. Papadimitriou, and Scott Shenker

Context

- The internet has over 12,000 autonomous systems
- Everyone picks their own upstream and/or peers
- AT&T wants to be close to everyone else on the network, but doesn't care about the network at large

Question:

What is the “penalty” in terms of poor network structure incurred by having the “users” create the network, without centralized control?

In this talk we…

- Introduce a simple model of network creation by self-interested agents
- Briefly review game-theoretic concepts
- Talk about related work
- Show bounds on the “price of anarchy” in the model
- Discuss extensions and open problems we believe to be relevant and potentially tractable.

A Simple Model

N agents, each can buy (undirected) links to a set of others (s_i)

One agent buys a link, but anyone can use it

Undirected graph G is built

Cost to agent:

$$c(i) = \alpha \cdot |s_i| + \sum_j d_G(i,j)$$

Pay $α for each link you buy
(α may depend on n)

Pay $1 for every hop to every node

Example

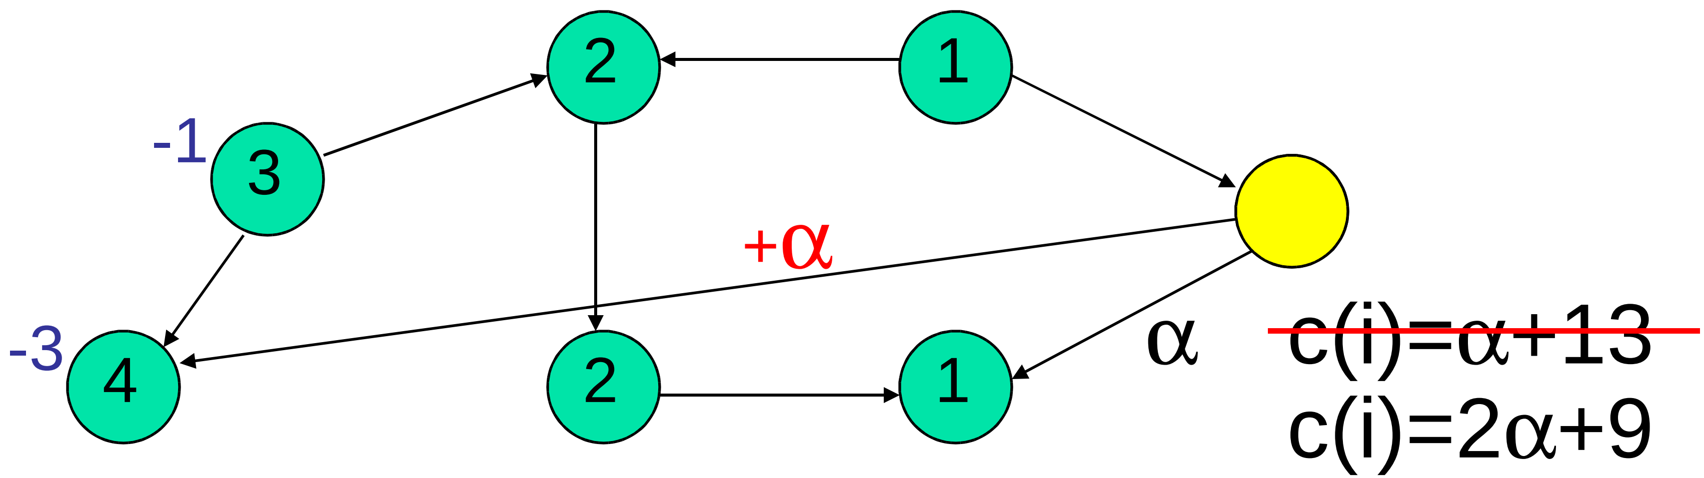

(Convention: arrow *from* the node buying the link)

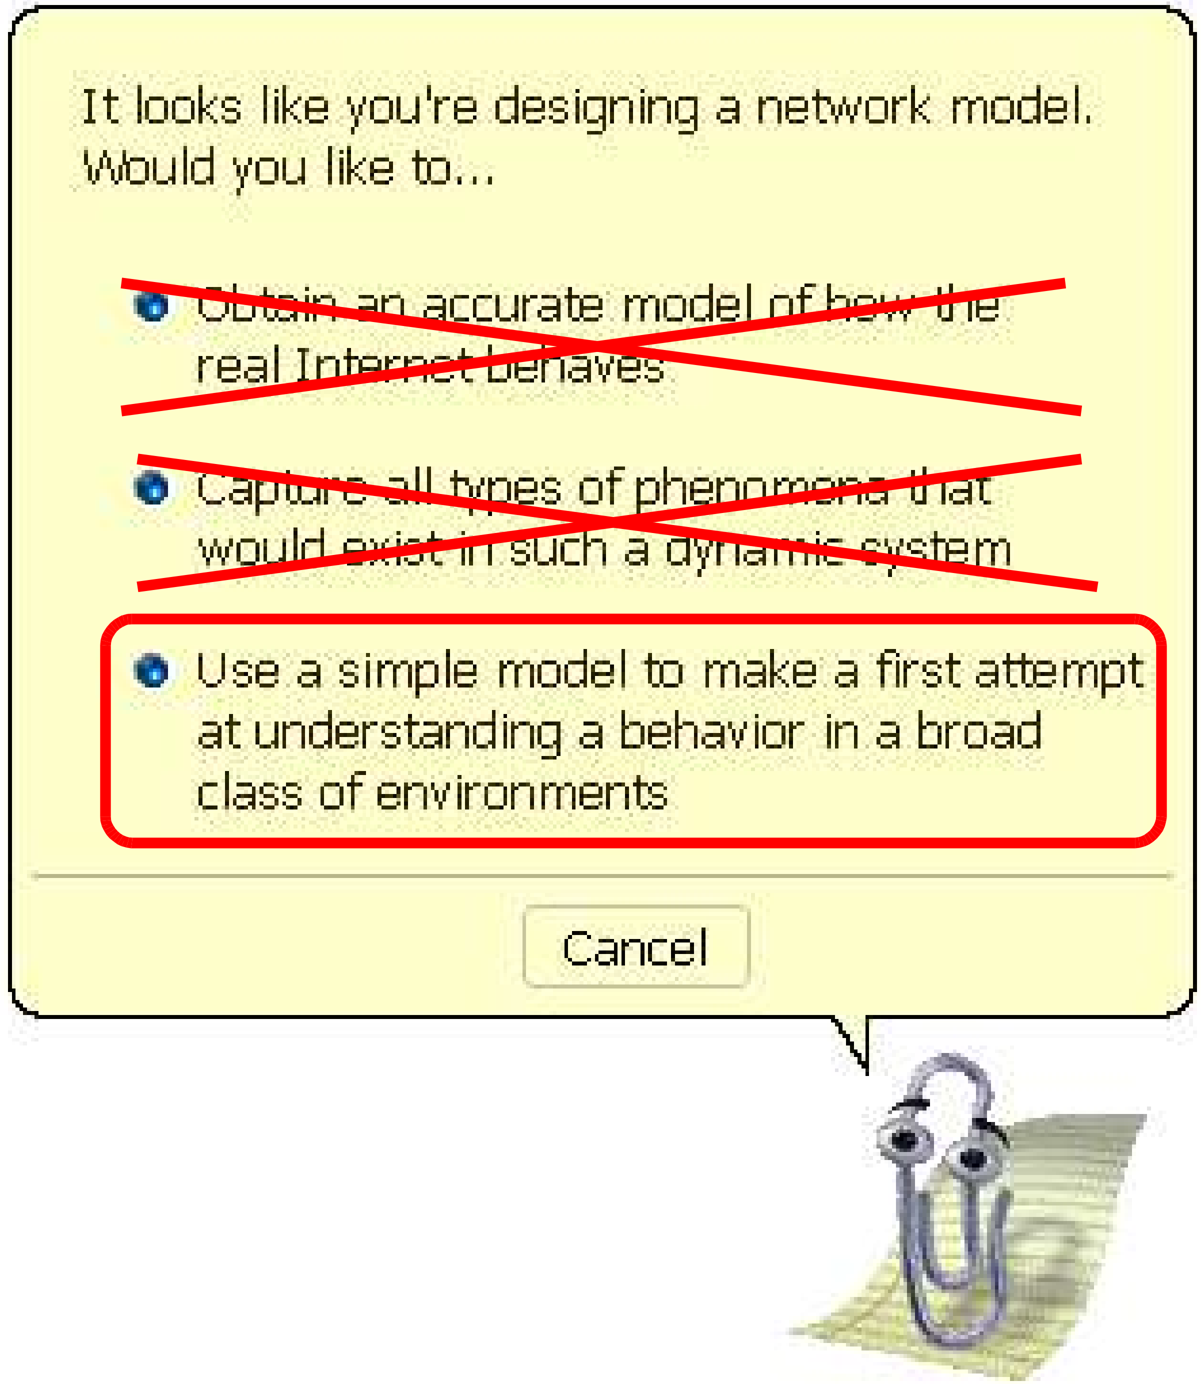
It looks like you're designing a network model. Would you like to...
Obtain an accurate model of how the real Internet behaves
Capture all types of phenomena that would exist in such a dynamic system
Use a simple model to make a first attempt at understanding a behavior in a broad class of environments
Cancel

Definitions

- Social cost: $C(G) = \sum_i c(i)$
  - The simplest notion of “global benefit”
- Social optimum: combination of strategies that minimizes the social cost
  - “What a benevolent dictator would do”
  - Not necessarily palatable to any given agent

Definitions: Nash Equilibria

- Nash equilibrium: a situation such that no single player can change what he is doing and benefit
  - A well-studied notion of “stability” in games, but not uncontroversial
  - Presumes complete rationality and knowledge on behalf of each agent
  - Not guaranteed to exist, but they do for our model

Example!

Set α=5, and consider:

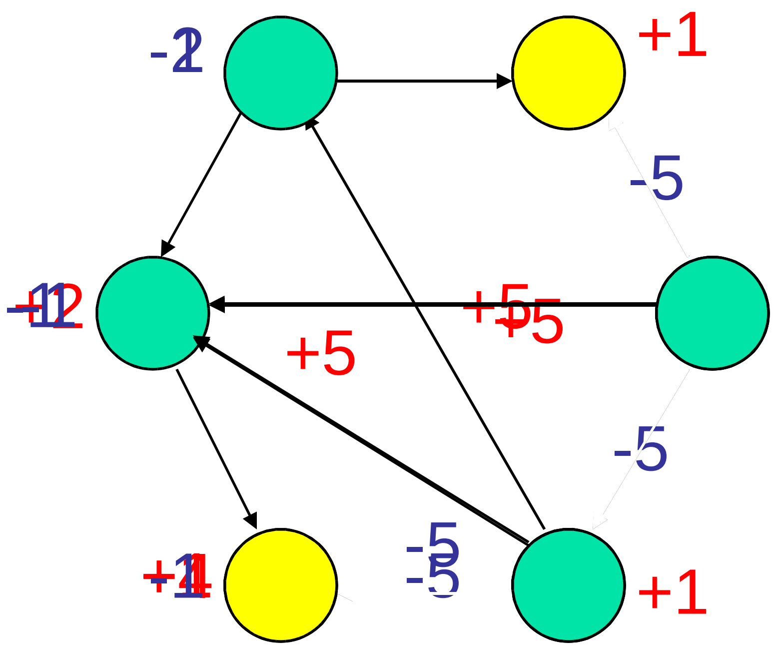

Definitions: Price of Anarchy

- Price of Anarchy (Koutsoupias & Papadimitriou, 1999): the ratio between the worst-case social cost of a Nash equilibrium network and the optimum network
- We bound the worst-case price of anarchy to evaluate “the price we pay” for operating without centralized control

Related Work

- Anshelevich, et al. (STOC 2003)
  - Agents are “global” and pick from a set of links to connect between their own terminals
  - Results concern the “optimistic price of anarchy” (with best-case Nash equilibria)
- A body of similar work on social networks in the econometrics literature (e.g. Bala&Goyal 2000, Dutta&Jackson 2000)

Our Results

- Complete characterization of the social optima
- Lower and upper bounds on the price of anarchy, constant in n, not tight in α
- A tight upper bound contingent on an experimentally-supported conjecture

Social optima

When $\alpha<2$, any missing edge can be added at cost α and subtract at least 2 from social cost

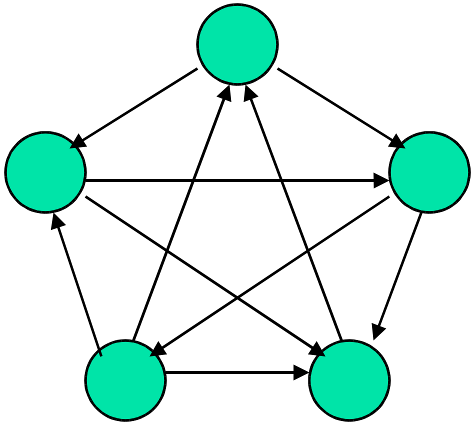

When $\alpha \geq 2$, consider a star. Any extra edges are too expensive.

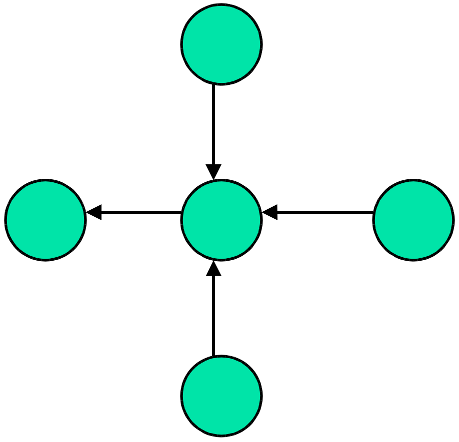

Equilibria: very small α (<2)

For α<1, the clique is the only N.E.

For 1<α<2, clique no longer N.E., but the diameter is at most 2; else:

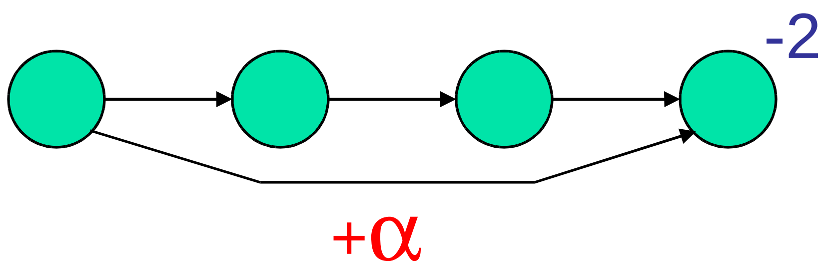

Then, the star is the worst N.E., can be seen to yield P.o.A. of at most 4/3

General Upper Bound

Assume α>2 (the interesting case)

Lemma: if G is a N.E., $d_G(i,j) < 2\sqrt{\alpha}$

Generalization of the above:

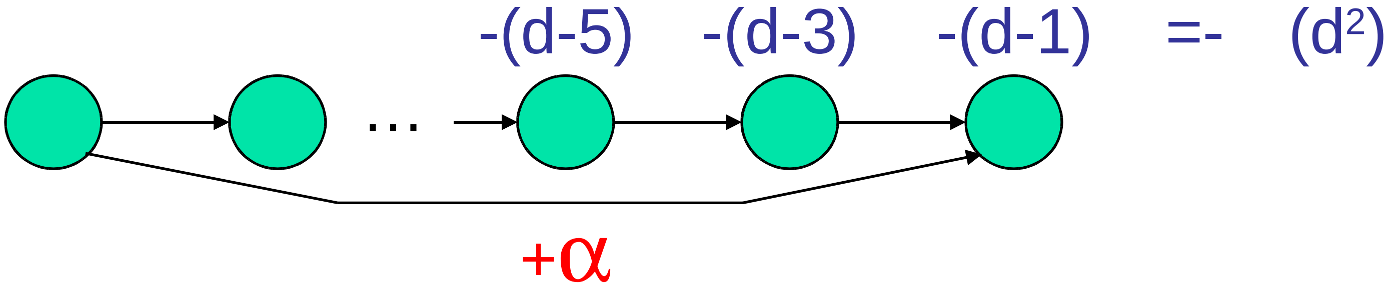

General Upper Bound (cont.)

A counting argument then shows that for every edge present in a Nash equilibrium, ($\sqrt{\alpha}$) others are absent

Then:

$$\begin{aligned} C(G) &\le \alpha \cdot O(n^2/\sqrt{\alpha}) + n(n-1) \cdot 2\sqrt{\alpha} \\ &= O(\sqrt{\alpha} n^2) \end{aligned}$$

C(star)= (n^2), thus P.o.A. is O($\sqrt{\alpha}$)

A Lower Bound

An outward-directed complete k-ary tree of depth d, at α=(d-1)n:

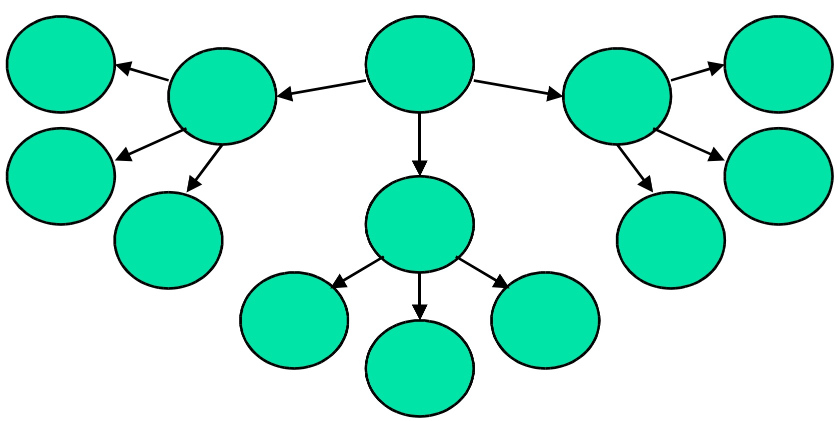

Infinite penalty for dropping existing links

No new link can bring you more than (d-1) closer to other nodes on *average*

A Lower Bound

An outward-directed complete k-ary tree of depth d, at $\alpha=(d-1)n$:

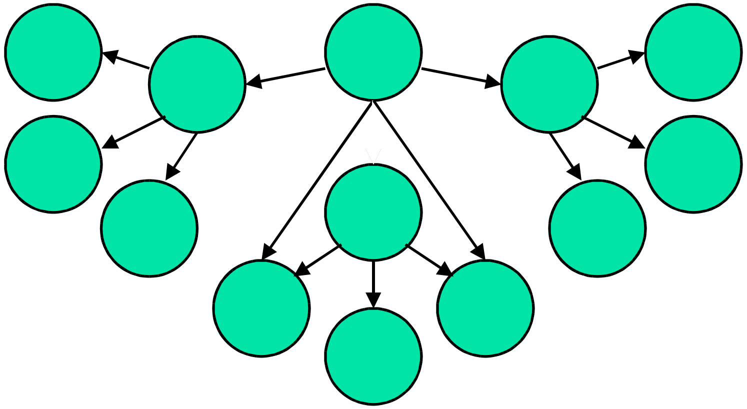

Can’t benefit from moving your existing links (the center of each subtree is the optimal site to link to)

A Lower Bound

An outward-directed complete k-ary tree of depth d, at $\alpha=(d-1)n$:

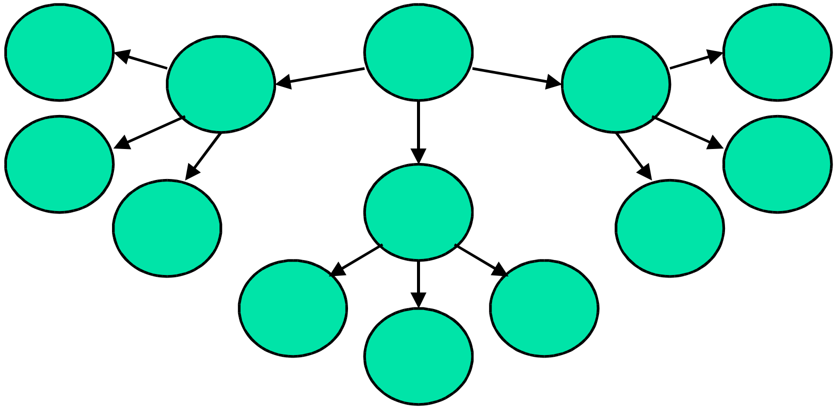

Benefit from adding several links is convex (net gain $\leq \Sigma$ individual gains), so won't create several new ones either

A Lower Bound

An outward-directed complete k-ary tree of depth d, at $\alpha=(d-1)n$:

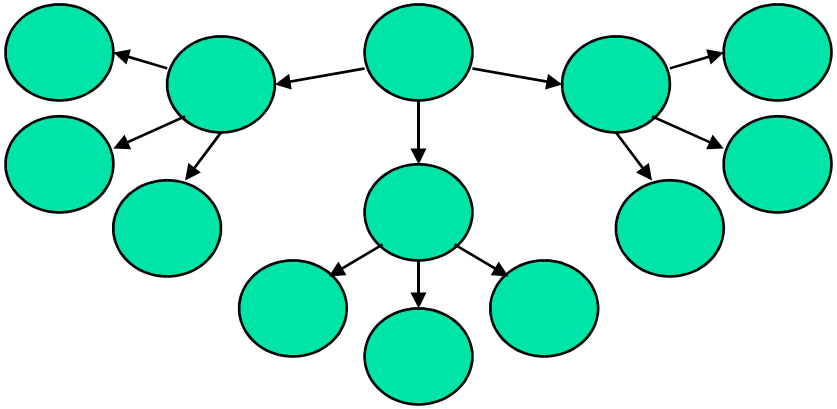

For large d, k, the price of anarchy approaches 3 asymptotically, so $3-\varepsilon$ is a lower bound for any $\varepsilon>0$

So what sorts of equilibria do exist?

Experimental Approach 1

- "Simulation":
  - Take a random ($G_{n,p}$) graph, iteratively have each agent re-optimize strategy until stable
  - But_1: no guarantee of convergence (although converges in practice)
  - But_2: each iteration is coNP-hard (simple reduction from Dominating Set)
- For $\alpha>2$, only trees observed, most often stars

Experimental Approach 2

- Application of the Feynman Problem-Solving Algorithm:
  - Write down n
  - Think really hard
  - Write down a non-tree Nash equilibrium
- Third step consistently fails
- Sole exception: the Petersen graph for $\alpha<4$, but still transient

Trees

- Conjecture: for $\alpha > \alpha_0$, some constant, all Nash equilibria are trees
- Benefit: a tree has a center (a node that, when removed, yields no components with more than n/2 nodes)
- Given a tree N.E., can use the fact that no additional nodes want to link to center to bound the depth and show that the price of anarchy is **at most 5**

Discussion

- The price tag of decentralization in network design appears modest
  - not directly dependent on the size of the network being built
- The Internet is not strictly a clique, or a star, or a tree, but often resembles one of these at any given scale
- Many possible extensions remain to be explored

Some Possible Extensions

- What if agents collaborate to create a link?
  - Each node can pay for a fraction of a link; link built only if total “investment” is ≥ 1
  - May yield a wider variety of equilibria
- Stars are efficient for hop distances, but problematic for congestion
  - What happens when agent costs are penalized for easily-congestible networks?

Some More Possible Extensions

- Most agents don't care to connect closely to everyone else
  - What if we know the amount of traffic between each pair of nodes and weight the distance terms accordingly?
  - If n^2 parameters is too much, what about restricted traffic matrices?
- Prevent "perfect blackmail" by making the penalty for complete disconnection large but finite?

Even More Possible Extensions

- Charge nodes for Vickrey payments? (from FPSS 2002)
- Introduce time?
  - Network develops in stages as new nodes arrive
  - Assume equilibrium state is reached at every stage

Directions for Future Work

- Proof of tree conjecture
- Price of anarchy in the above models
- Other points on the spectrum between dictatorship and anarchy?
- Measurements to assess applicability to existing real systems

Questions?